Science at the Petascale

Douglas B. Kothe

Oak Ridge Leadership Computing Facility

Presentation to the Cray XT5 Quad-Core Workshop

April 13, 2009

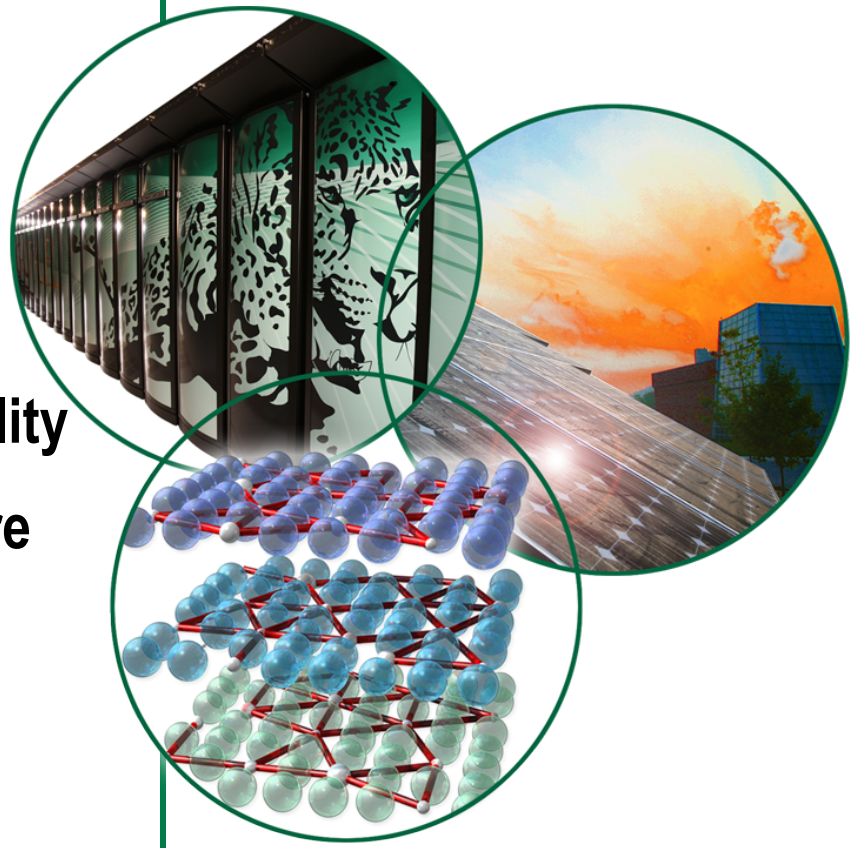

U.S. DEPARTMENT OF ENERGY

OAK RIDGE NATIONAL LABORATORY

MANAGED BY UT-BATTELLE FOR THE DEPARTMENT OF ENERGY

ORNL is the U.S. Department of Energy's largest science and energy laboratory

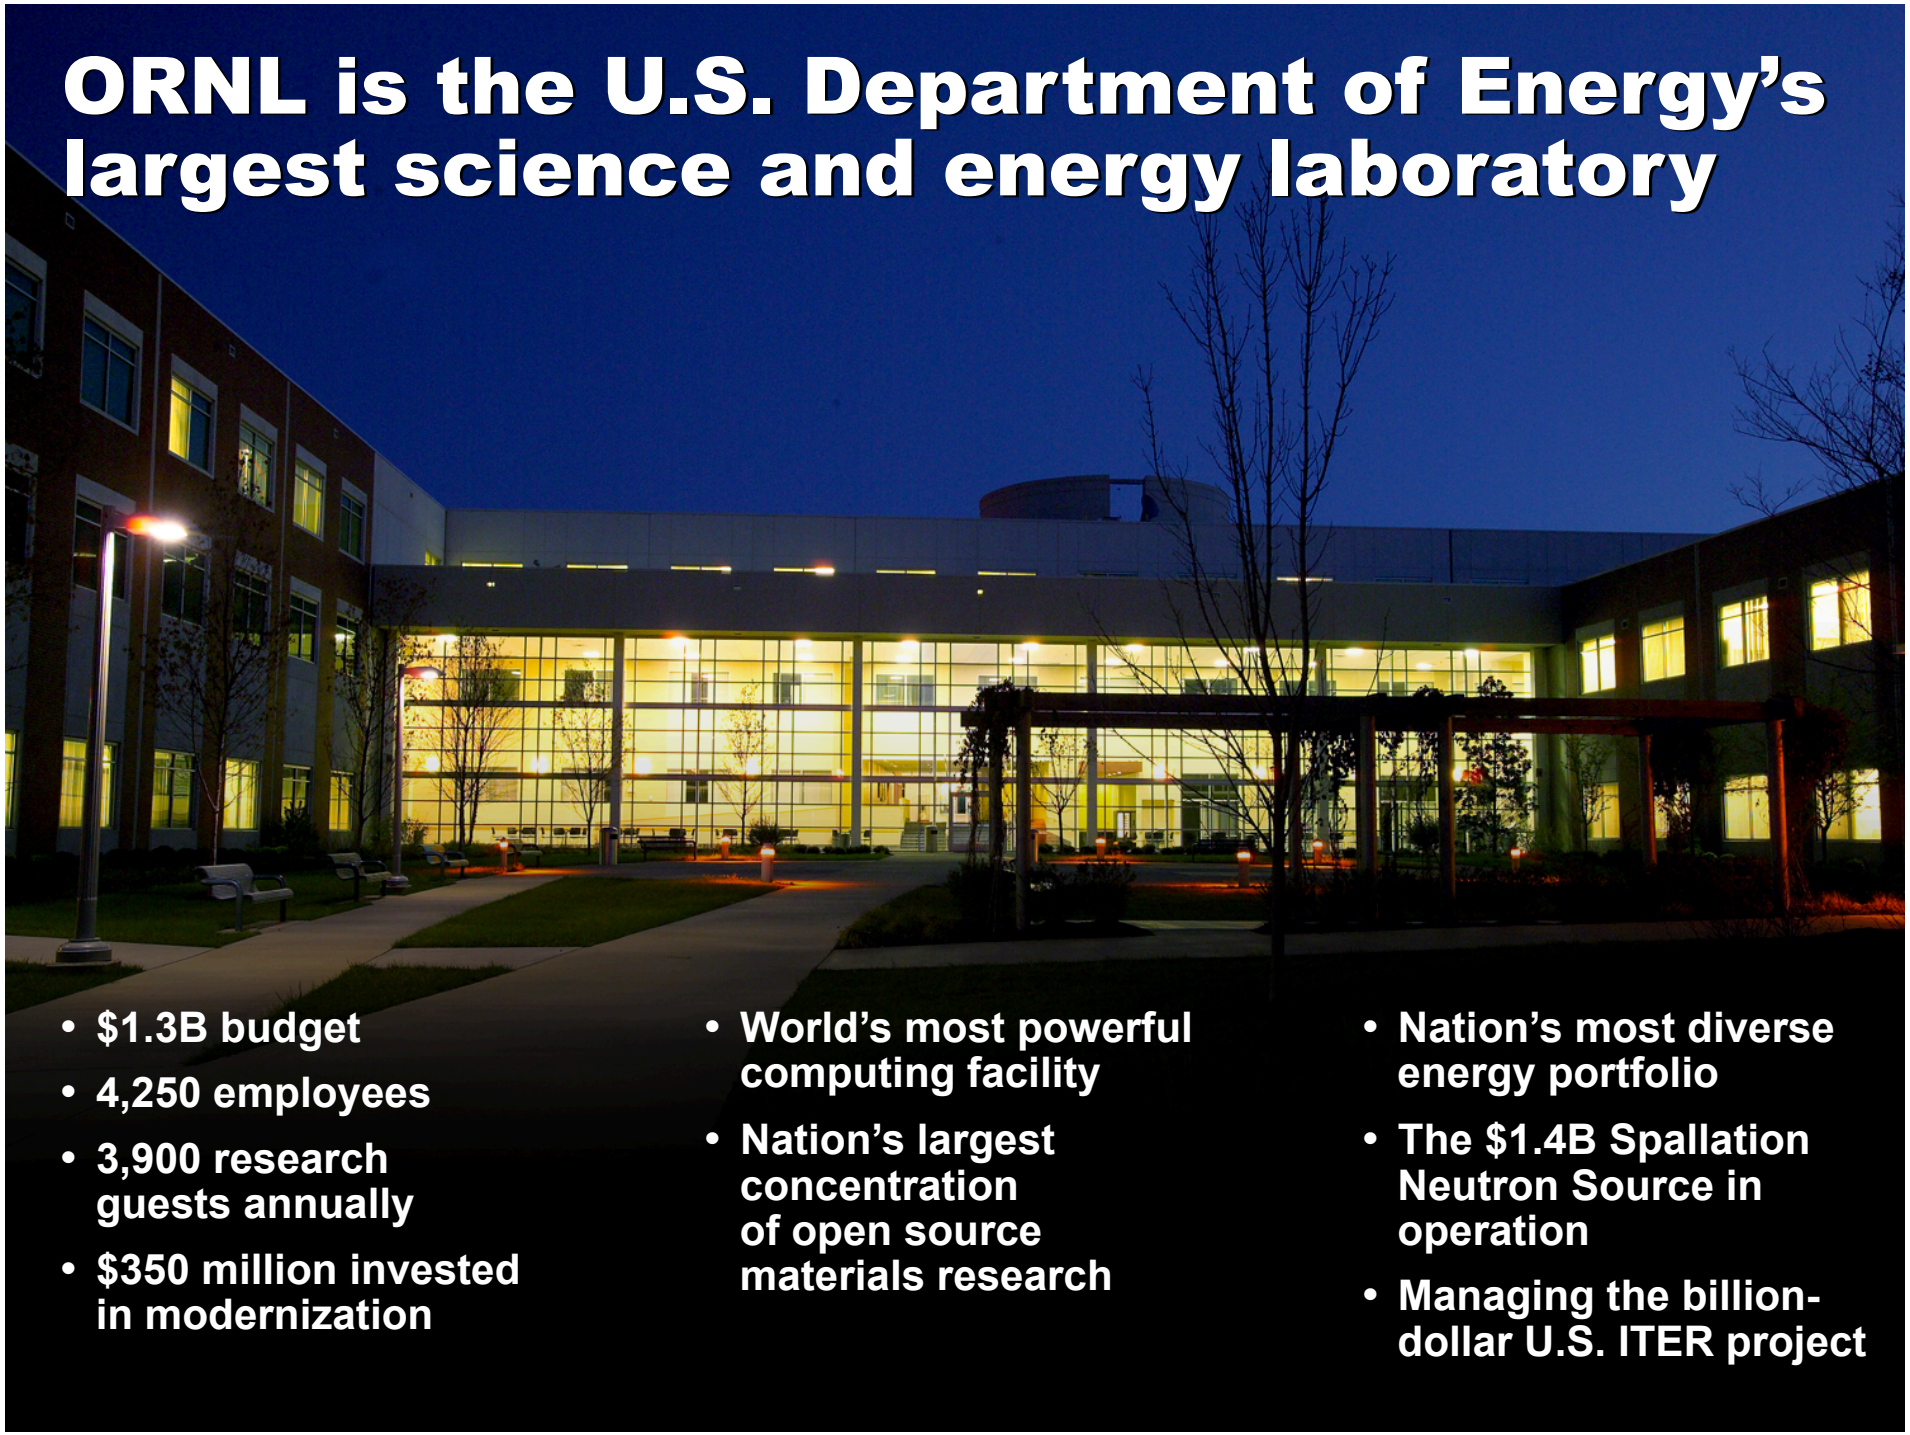

- $1.3B budget
- 4,250 employees
- 3,900 research guests annually
- $350 million invested in modernization

- World's most powerful computing facility
- Nation's largest concentration of open source materials research

- Nation's most diverse energy portfolio
- The $1.4B Spallation Neutron Source in operation
- Managing the billion-dollar U.S. ITER project

National user facilities at ORNL

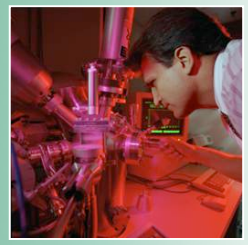

- Center for Nanophase Materials Sciences
- Cooling, Heating, and Power Integration Laboratory
- Fuels, Engines, and Emissions Research Center
- High Flux Isotope Reactor
- High Temperature Materials Laboratory
- Holifield Radioactive Ion Beam Facility
- Metals Processing Laboratory Users Facility
- Mouse Genetics Research Facility
- National Transportation Research Center
- Oak Ridge Electron Linear Accelerator
- Spallation Neutron Source

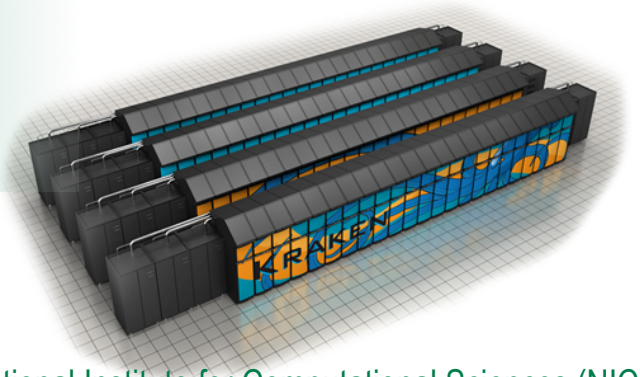

National Institute for Computational Sciences (NICS)

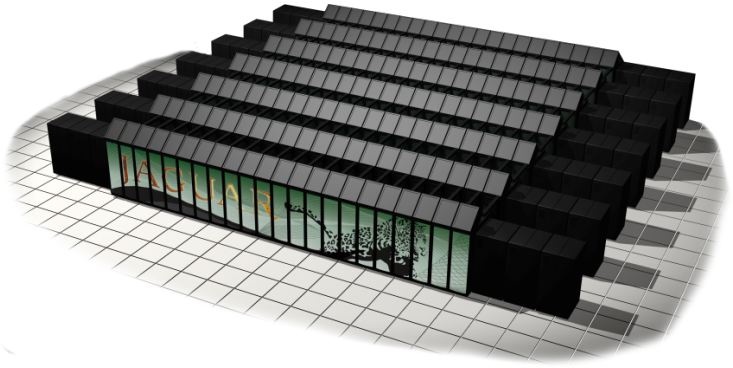

National Center for Computational Sciences (NCCS)

We are the nation's lead laboratory for scientific computing

NCCS
National Center for Computational Sciences

NICS
National Institute for Computational Sciences

Applications

Superconductivity | Chemistry | Fission | Climate change | Bioinformatics
Materials | Combustion | Fusion | Geosciences | Biology
Nanoscience | Energy | Molecular dynamics | Groundwater | Biophysics

Institute for Advanced Architecture and Algorithms
DOE-SC and NNSA

Extreme Scale Software Center
DOE-SC and DOD

HPC experience and operations

Facilities and infrastructure

Data storage and file systems

Computer systems

Visualization

Networks

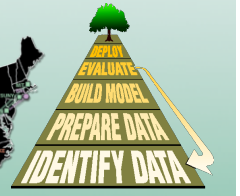

Knowledge discovery

OAK RIDGE National Laboratory

Science advances for the next decade require leadership computing

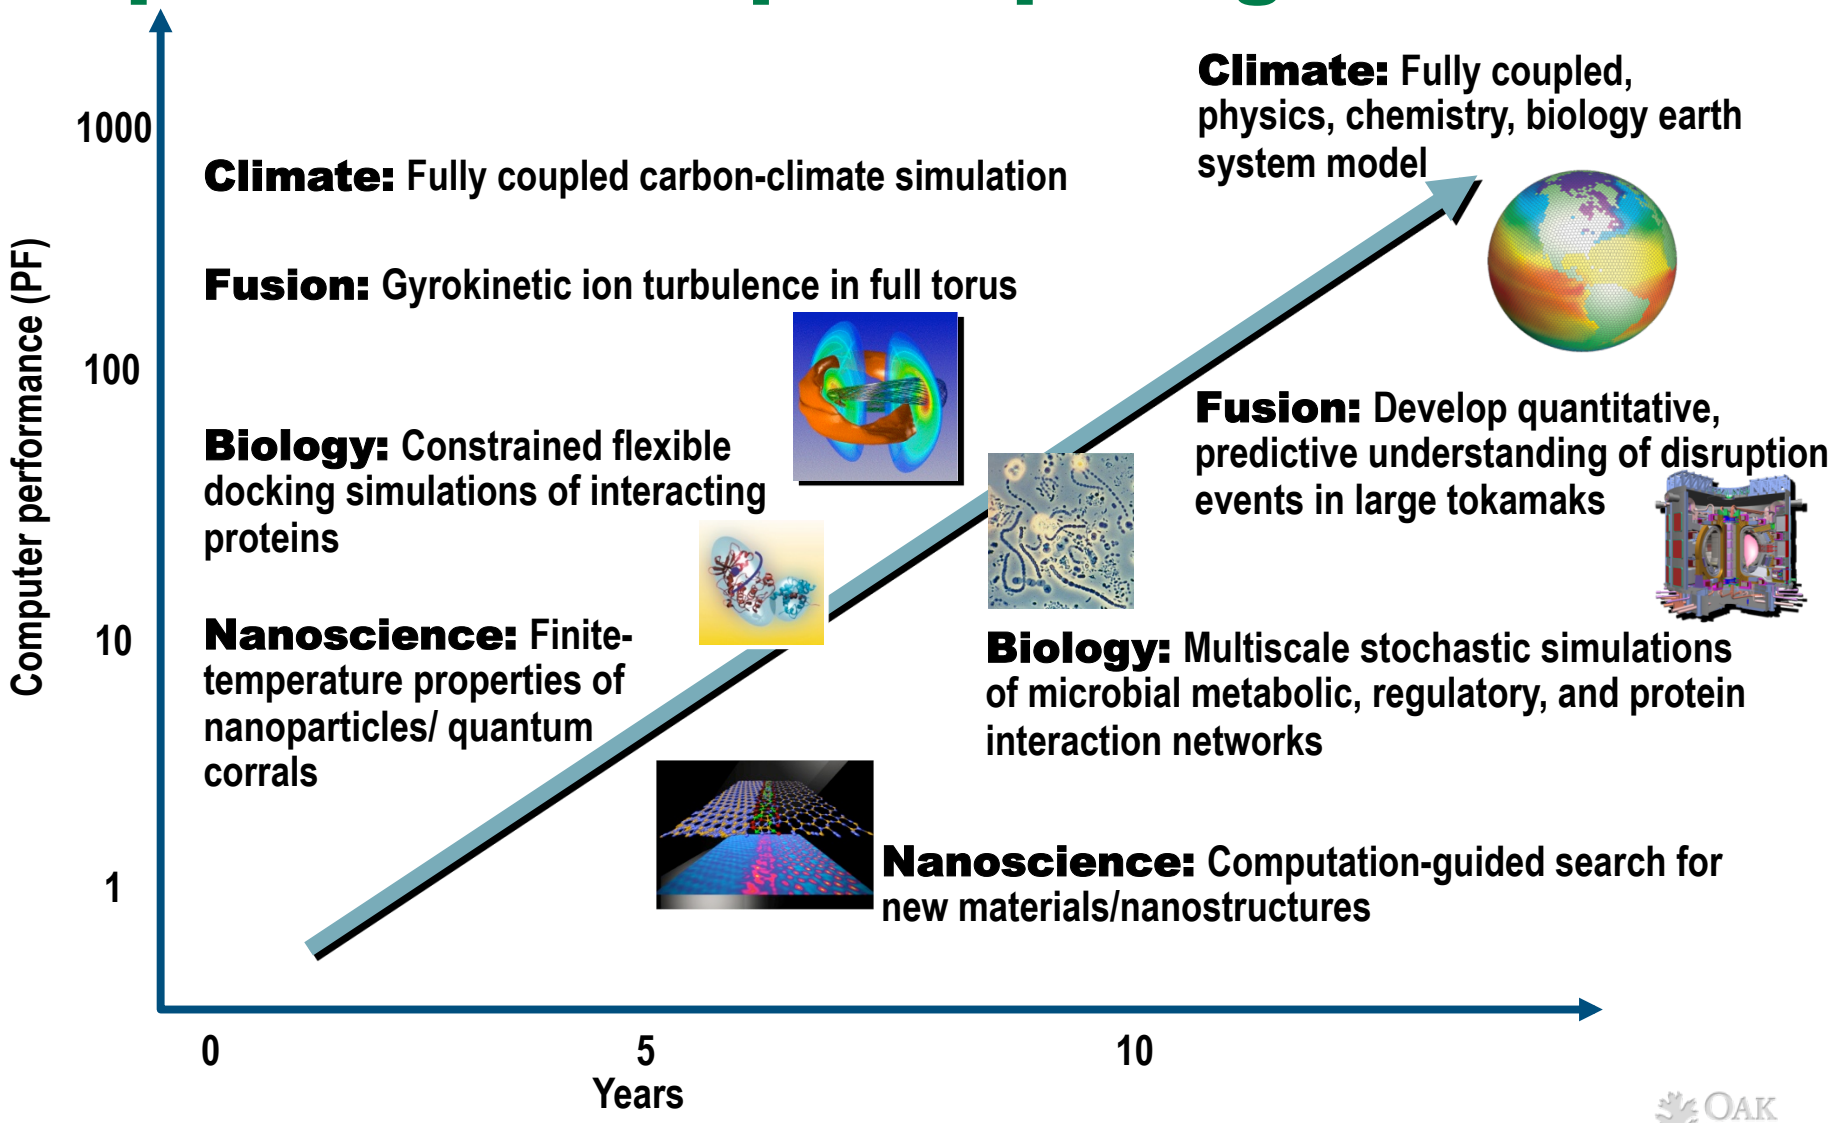

National Center for Computational Sciences
Oak Ridge National Laboratory

Mission: Deploy and operate the computational resources required to tackle global challenges

- Providing world-leading computational resources and specialized services for the most computationally intensive problems
- Providing stable hardware/software path of increasing scale maximize productive applications development
- Deliver transforming discoveries in materials, biology, climat energy technologies, etc.
- Ability to investigate otherwise inaccessible systems, from supernovae to energy grid dynamics

Jaguar – 1.64 PF Cray XT: 45,376 Quad-Core Processors, 362 TB memory

Jaguar: World's most powerful computer Designed for science from the ground up

Peak performance	1.645 petaflops
System memory	362 terabytes
Disk space	10.7 petabytes
Disk bandwidth	200+ gigabytes/second

DOE's broad range of science challenges demand a balanced, scalable, general purpose system

Projects	2006	2007	2008	2009
Accelerator physics	1	1	1	1
Astrophysics	3	4	5	5
Chemistry	1	1	2	4
Climate change	3	3	4	5
Combustion	1	1	2	2
Computer science	1	1	1	1
Fluid Dynamics			1	1
Fusion	4	5	3	5
Geosciences		1	1	1
High energy physics		1	1	
Life sciences	2	2	2	4
Materials science	2	3	3	4
Nuclear physics	2	2	1	2
Industry	2	3	3	3
Total Projects:	22	28	30	38
CPU Hours (Millions):	36	75.5	145	470

National Institute for Computational Sciences

A UT/ORNL PARTNERSHIP
NATIONAL INSTITUTE FOR COMPUTATIONAL SCIENCES
NICS

- Awarded a $65M NSF grant to deploy and operate the world's most powerful academic supercomputer for the US research community
- NICS is a collaboration between the University of Tennessee and ORNL with facilities located on the ORNL campus

Kraken Upgrade World's Most Powerful Academic Supercomputer

A UT/ORNL PARTNERSHIP
NATIONAL INSTITUTE FOR COMPUTATIONAL SCIENCES
NICS

Cray XT5 system – Kraken

- 88 cabinets
- 8,256 dual-socket nodes
- 16,512 quad-core Opterons
- 66,048 cores
- 100 TB memory
- SeaStar2+ interconnect
- 30 GB/s disk bandwidth
- 3.3 PB disk (2.6 formatted)
- 600+ peak TFLOPS
- Plan to upgrade to 6-core chips in 2009

Jaguar's Cray XT5 Nodes Designed for science

- **Powerful node improves scalability**
- **Large shared memory**
- **OpenMP Support**
- **Low latency, High bandwidth interconnect**
- **Upgradable processor, memory, and interconnect**

GFLOPS	76.3
Memory (GB)	16
Cores	8
SeaStar2+	1

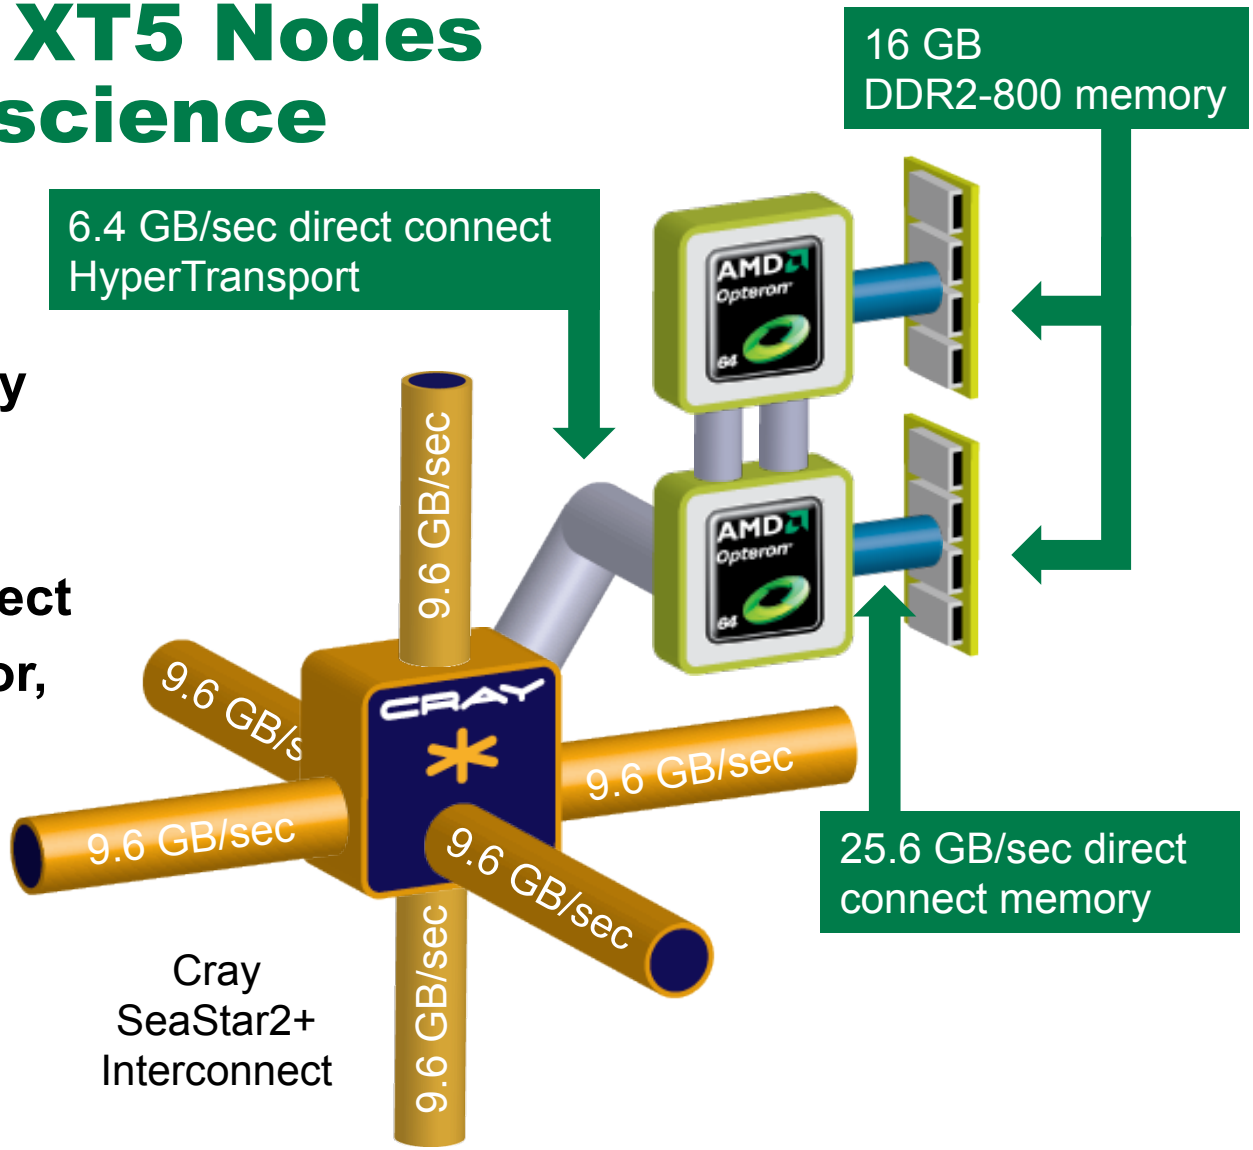

Storage for an avalanche of data

Center wide shared file systems

- **"Spider" will be available later this year to provide a shared, parallel file system for all LCF systems**
  - **Based on Lustre file system**
- **Planned bandwidth of over 200 GB/s with multi-petabytes (10 PB) of capacity**
- **HPSS provides archival storage for all system**
- **HPSS has been upgraded with two additional tape libraries to add additional capacity and bandwidth**

Visualization and Data Analytics

Visualization

Once users have completed their runs, the Visualization task group helps them make sense of the sometimes overwhelming amount of information they generate.

- Viewing at a 30'x8' PowerWall
- Upgraded cluster with GPUs for remote visualization

End-to-End Solutions

Researchers must analyze, organize, and transfer an enormous quantity of data. The End-to-End task group streamlines the work flow for system users so that their time is not eaten up by slow and repetitive chores.

- Automate routine activities, ex. job monitoring at multiple sites
- Data Analysis

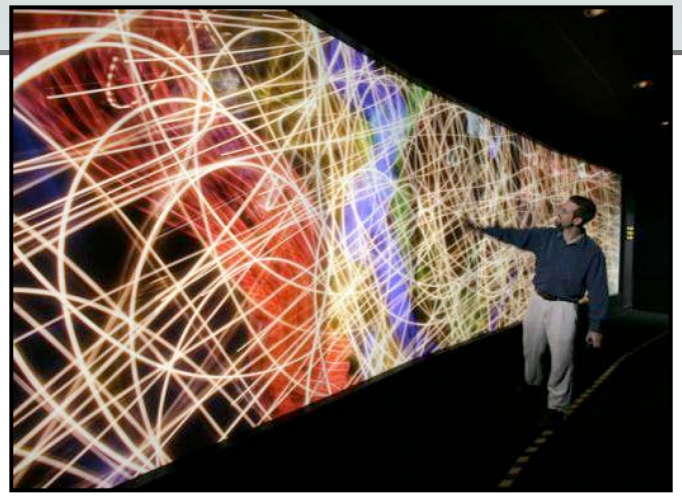

Gordon Bell prize awarded to ORNL team

Three of six GB finalist ran on Jaguar

- A team led by ORNL's Thomas Schulthess received the prestigious 2008 Association for Computing Machinery (ACM) Gordon Bell Prize at SC08
- For attaining fastest performance ever in a scientific supercomputing application
- Simulation of superconductors achieved 1.352 petaflops on ORNL's Cray XT Jaguar supercomputer
- By modifying the algorithms and software design of the DCA++ code, the team was able to boost its performance tenfold

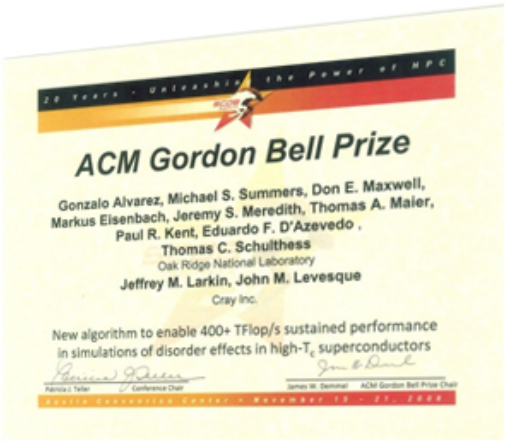
ACM Gordon Bell Prize

Gonzalo Alvarez, Michael S. Summers, Don E. Maxwell, Markus Eisenbach, Jeremy S. Meredith, Thomas A. Maier, Paul R. Kent, Eduardo F. D'Azevedo , Thomas C. Schulthess
Oak Ridge National Laboratory
Jeffrey M. Larkin, John M. Levesque
Cray Inc.

New algorithm to enable 400+ TFlop/s sustained performance in simulations of disorder effects in high-T_c superconductors

Gordon Bell Finalists	
✓ DCA++	ORNL
✓ LS3DF	LBNL
✓ SPECFEM3D	SDSC
• RHEA	TACC
• SPaSM	LANL
• VPIC	LANL

OAK RIDGE National Laboratory

Pushing Back the Frontiers of Science

Petascale Early Science Projects Tackle National/Global Problems

- **Energy for environmental sustainability**
  - **Climate change:** carbon sequestration, weather event impacts on global climate, decadal climate predictive skill in aerosol forcing, global climate at unprecedented resolution
  - **Bioenergy:** recalcitrance in cellulosic ethanol
  - **Solar:** non-equilibrium semiconductor alloys
  - **Energy storage:** charge storage and transfer in nano-structured supercapacitors
  - **Energy transmission:** role of inhomogeneities in high-T superconducting cuprates
  - **Combustion:** stabilizing diesel jet flames for increased efficiency & decreased emissions
  - **Fusion:** ITER design, optimization, and operation
  - **Nuclear energy:** fully-resolved reactor core neutron state
- **Materials and nanoscience**
  - **Structure** of nanowires, nanorods, & strongly correlated materials (magnets)
- **Fundamental science**
  - **Astrophysics:** decipher core-collapse supernovae & black hole mergers
  - **Chemistry:** elucidate water structure in biological & aqueous-phase systems
  - **Nuclear physics:** probe the anomalously long lifetime of Carbon-14
  - **Turbulence:** dispersion relative to air quality modeling and bioterrorism

Million-fold increase in computing and data capabilities

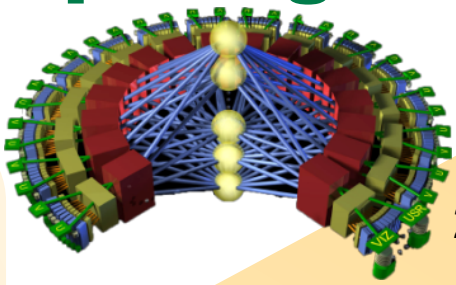

2004
Cray X1
3 TF

2005
Cray XT3
Single-core
26 TF

2006
Cray XT3
Dual-core
54 TF

2007
Cray XT4
119 TF

2008
Cray XT4
Quad-core
166 TF
18TB, 0.3 PB

2009
Cray XT4
Quad-core
263 TF
62 TB, 1 PB

Jaguar
8-core, dual-socket SMP
1.64PF
362 TB, 10 PB

Cray XT5
8/12-core, dual-socket SMP
~1 PF
100TB, 3.3 PB

2011
DARPA HPCS
20 PF

2015
Future system
100–250 PF

2018
Future system
1000 PF
(1 EF)

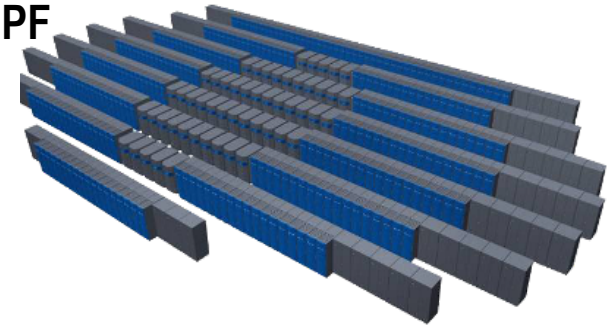

National HPC centers: It's the people!

Unsurpassed in-house expertise working with users to effectively use computational science, visualization, and workflow technologies on HPC resources to

Port, tune, augment, and develop current and future applications at scale

Easily access HPC systems

Provide visualizations to present scientific results and augment discovery processes

Automate the scientific computational method

Communicate your accomplishments to a broad audience